N91-20686

The Real-Time Control of Planetary Rovers Through Behavior Modification

David P. Miller
Jet Propulsion Laboratory/California Institute of Technology
M.S. 301-440, 4800 Oak Grove Drive
Pasadena, CA 91109
(818) 354-9390
dmiller@ai.nasa.jpl.nasa.gov

Abstract

It is not yet clear of what type, and how much, "intelligence" is needed for a planetary rover to function semi-autonomously on a planetary surface. Current designs assume an advanced AI system that maintains a detailed map of its journeys and the surroundings, and that carefully calculates and tests every move in advance. To achieve these abilities, and because of the limitations of space-qualified electronics, the supporting rover is quite sizable, massing a large fraction of a ton, and requiring technology advances in everything from power to ground operations.

An alternative approach is to use a behavior driven control scheme. Recent research has shown that many complex tasks may be achieved by programming a robot with a set of behaviors and activating or deactivating a subset of those behaviors as required by the specific situation in which the robot finds itself. Behavior control requires much less computation than is required by traditional AI planning techniques. The reduced computation requirements allows the entire rover to be scaled down as appropriate (only down-link communications and payload do not scale under these circumstances). This paper discusses the missions that can be handled by the real-time control and operation of a set of small, semi-autonomous, interacting, behavior-controlled planetary rovers.

1. Introduction:

There are many possible uses for unmanned planetary rovers. Rovers with a high degree of autonomy can carry out missions that serve science, operations, and space exploitation goals. For example, rovers can be used on the Moon to perform site certification for possible manned outposts and science instrument sites. On Mars, science instruments need to be placed and soil and rock samples need to be gathered from a wide variety of terrains. To reduce light-time delays and the need for communications (and its inherent infrastructure of relay satellites etc), rovers with at least semi-autonomous capabilities, are highly desired.

1.1 Plan Control for Rovers

The autonomous system control that has been proposed for a rover, to accomplish the tasks mentioned above, is shown in Figure 1. The rover senses its environment, combines that with previous knowledge (from earlier and orbital views) and then builds a map of its surroundings. A path planner finds a trajectory through the map. The trajectory is simulated, producing run-time expectations, and these expectations are monitored during the actual rover traverse. If an expectation is violated, the rover performs a reflex stop, and starts the cycle over again. Under normal circumstances the cycle is repeated every five to fifteen meters. Such a system has been successfully implemented, and tested under realistic conditions [Miller89, Gat90].

The implemented system required just under one billion machine instructions per meter of travel. While the code used in this experimental system was by no means optimal, by the time all the functionality and reliability improvements are made, it is believed that the real number will be within a factor of three. This means that a rover that needed to travel at a speed of one kilometer per hour would need a (space qualified) computational capability of between 80 and 250 MIPS. Some of this could be offloaded onto special purpose computation systems, but none the less, computation becomes a major driver for a planetary rover, both in power and mass.

The system studies that have been undertaken [Pivirotto90] have confirmed this, indicating that a rover with some onboard autonomy would need to mass between 600 and 842kg (the test vehicle we used massed over 1100kg). In large part, this mass is due to the system control algorithms.

1.2 Low-Mass Rovers are Needed

For almost all imagined uses of rovers, the job is best done if it can be done several times, in many places, under differing conditions. Traditional rovers weigh several hundred kilograms, and under the best circumstances will probably be able to move a few kilometers a day. It will not be economically feasible to place more than a few such rovers on a planetary surface. Nor will it be logistically possible to have those few rovers visit all the desired sites. Additionally, the risk of losing one of these scarce and expensive resources will make it difficult to put a rover in the places where it could do the most good: in the previously unmapped or geologically unknown areas of a planetary surface.

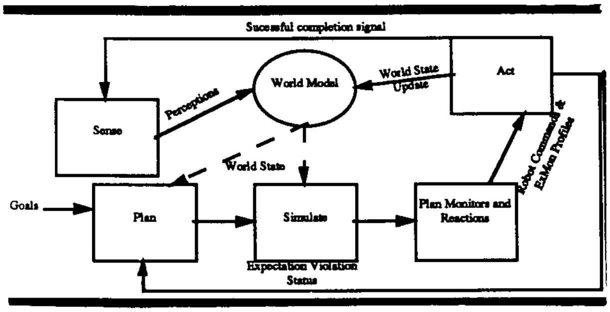

Figure 1. Planning Cycle

All of the problems mentioned above go away if a low-computation method for autonomously controlling rovers could be developed. Low computation greatly lowers the power needs of the rover, which can greatly reduce its mass. Once the power mass cycle is broken, it is possible to scale the entire rover. The strength to weight ratio of mechanical systems improves as a system is scaled smaller. This in turn makes the system stronger and stiffer, which will reduce the accuracy of the control needs, further reducing the power and payload needs of the rover, allowing it to be scaled even smaller. The rover cannot be made arbitrarily small, because it is necessary for it to carry some sort of payload, but improvements in instrument miniaturization [Waltman89] indicate that that may not be much of a limitation. A rover must also travel through its environment, but remember that an ant can travel anywhere a person can, and many more places in addition.

To keep this all within todays technology, if one could build mini-rovers that massed a few kilograms, and could operate autonomously, then many useful things could be done with them. Mini-rovers could be sent in far greater numbers, and to potentially far riskier sites. It now appears feasible to design a small rover that can accomplish most of the science and operations objectives that previously would have been handled by a large rover. Key in building a mini-rover is behavior programming.

2. Behavior Control for Rovers

Behavior programming is a methodology that allows a set of interacting reactions to be programmed into a robot so that they work together. For example, if a robot has a behavior that causes the robot to turn so as to maximize the value being returned by its right side range finder, and has a behavior to cause the robot to turn away when its left side proximity sensor is activated, then that robot will exhibit the behavior of traveling around the perimeter of a room in a clockwise direction.

Some behavior languages have been created [Brooks86, Gat90], and several robots have been programmed with these languages to perform some interesting tasks [Brooks89, Gat90]. Central to having a robot perform something of interest is the ability to have the robot's behavior modified by cues in the environment, or by the robot's own actions.

Figure 2 shows the sensors and actuators of a small robot named "Tooth". Tooth is a robot massing just under two kilograms, and containing two eight-bit micro-processors and four kilobytes of memory. Its memory is filled with the behavior program shown in Figure 3. Tooth's mission is to go around a room, picking up "toys" that have been left near the walls of the room, and to bring the toys to a beacon located somewhere near the center of the room. The robot carries out this task using nine interacting behaviors. The look for toy behavior just has the vehicle move along in a straight line. The follow light behavior uses the photocells to decide on the direction of the beacon. If the beacon can be seen, this behavior sends a steering command away from the beacon. This forces the robot towards the walls. The dead-end, obstacle avoidance, unthrash, and stall behaviors all keep the robot from crashing into anything, or getting stuck. When the object beam sensor detects something, then the robot uses the pickup toy behavior to grasp the object. If it is successful, then that behavior causes the steering signal from the follow-light behavior to be inverted. This causes the robot to head towards the beacon. When the photocells report a light above a certain threshold brightness, then the drop toy behavior turns on. This overrides the look for toys behavior, causes the robot to stop, drop the toy and back away a few centimeters. Since the robot is no longer holding the toy, the follow-light signal is no longer inverted, and the robot turns away from the light and goes in search of more toys. [Gat90] gives more details on these experiments.

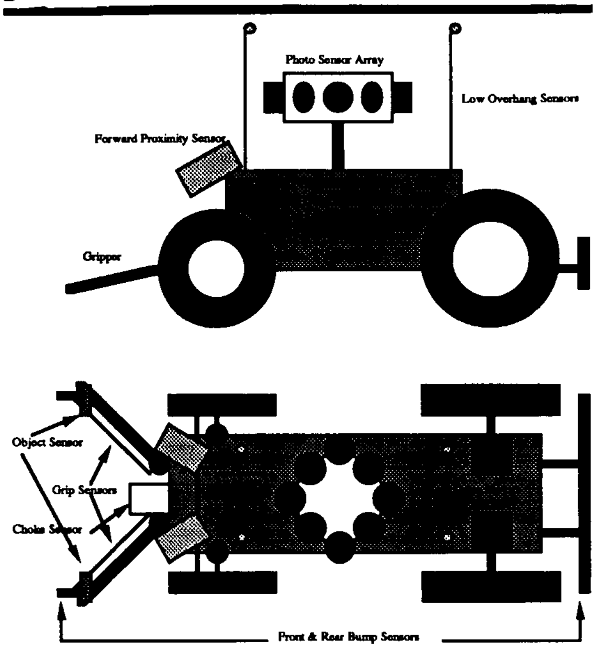

Figure 2: The Tooth Robot

The experiment outlined above shows that simple behaviors can be linked together to exhibit complicated actions. The experiment above had the robot perform all the major parts of a sample return mission. The robot moved following certain parameters. When it came across something that matched its sample criteria, it acquired it. It then brought the sample back to a marked beacon (simulating a return vehicle dock). All the while, the robot avoided obstacles in its path. This technology is sufficient for a variety of missions.

3. Target missions:

There are many missions that can be performed by behavior-controlled robots. The advantages of these small, autonomous rovers are many. Because of their small size, low mass, and high strength, these rovers could be placed on a planetary surface without the expense and mass of a traditional soft lander. [Miller90] outlines how many mini-rovers could be landed on Mars using parachutes and areoshells, and how communications can be maintained through ground relays. On the Moon, mini-rovers could be landed using an updated version of the Ranger seismometer capsule [Ranger63]. Navigation can be handled by referencing the robot to a coded radio signal, such as that used in VOR aviation radios.

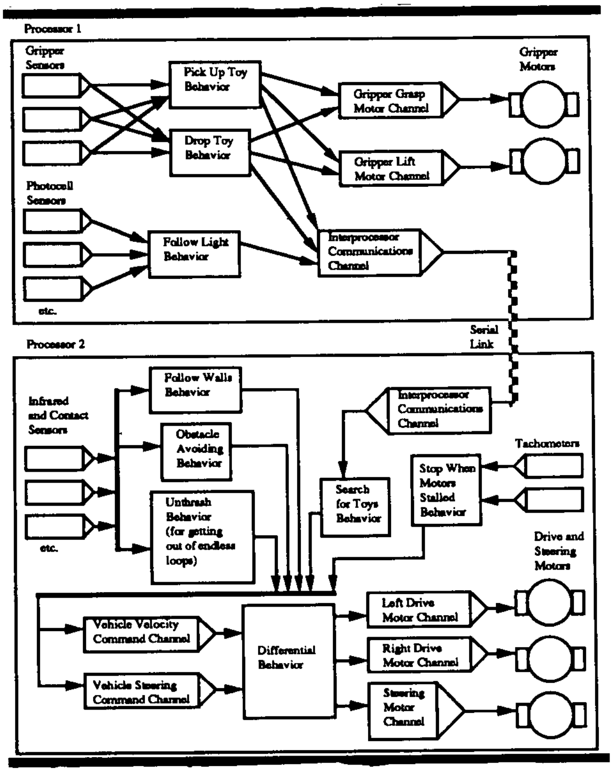

Figure 3: The Behavior Program for Tooth

3.1 A Lunar Mission

On the Moon, one of the major activities of rovers would be deploying science instruments (e.g., VLFA). The VLFA is large antenna consisting of several hundred elements laid out over a sixty kilometer long spiral arc. The exact placement is not crucial, but the elements should be distributed evenly along the arc. The elements themselves are self contained and mass approximately one kilogram. A previous study undertaken by the Battelle Corporation [Easter88] chose to use 1060kg rover that could carry the approximately 300kg of antenna elements. The rover contained large robotic arms for implanting the antenna elements. Since the VLFA is to be placed on the Lunar Farside, a series of relay satellites would be placed in lunar orbit to allow near continuous communications with the rover.

There are several ways the VLFA could be set up using behavior controlled mini-rovers. The simplest method would be to build one antenna element into each mini-rover, and land approximately 300 on the lunar farside. It would be easy to have the rovers distribute themselves relative to one another, in the desired pattern. This, most inefficient use of mini-rovers, would still result in a considerable mass savings over the Battelle proposal.

One could also send up, say fifty mini-rovers, that could each retrieve antenna elements from a central cache. It would then be trivial to have the rovers place the antenna elements in the proper arrangement about a VOR transmitter located at the cache. Each rover would have prestored the coordinates (radial and range) of the six antenna elements it was to place. Simple behaviors would cause the rovers to home in on the transmitter till they could spot an antenna element. Once they picked that up, they would circle the transmitter till they located the proper radial. They would then head outward along the radial the proper distance and deposit the antenna element. No longer holding anything, they would head back towards the cache and repeat the process with the next set of position coordinates. This scheme would require only about 500kg landed on the Moon, much lower lander masses, and only the relay satellites required by the VLFA.

Any activity that needs to take place in a pattern can easily be done with behavior-controlled mini-rovers. A simple non-directional radio-beacon along with a range encoder can be used to have a rover go in circles, spiral in or out, or rendezvous at the beacon. By adding the radial information that comes from a VOR, a mini-rover could have its behaviors direct it to a specific point relative to the beacon. All the while, other behaviors can be used to keep the rover from hanging up on obstacles, or getting stuck in dead-ends.

3.2 A Mars Mission

On Mars, there are several unanswered questions which require rovers to work at many diverse locations. In particular, the search for carbonates, and the emplacement of science stations (seismic and meteorological) are very suitable for behavior-controlled mini-rovers. These tasks require rovers in many different terrains in areas that are many thousands of kilometers apart, and are most useful in areas where the rover may never be able to get out (eg., inside extinct volcano craters). Here, the basic philosophy should be put down a lot of rovers, all over the place, and have them report back when they find something interesting.

More narrow scope missions on Mars are also applicable. A question, which could be key for the design of a sample return mission, is the thickness of the weathering rind on Mars rocks. Scientific opinion ranges from a millimeter to several centimeters. The extent of the rind will determine the type of coring that will be necessary for a sample return mission. A small rover, armed with a small circular saw, would be able to answer this question (or at least determine if the rind is more than a centimeter). Such a rover would have behaviors to get it close up to rocks of various sizes. It would have a list of criteria, as soon as it found a rock matching some of those, it would come up to the rock, and using the control behaviors, cut off a slice. It would then take several images and relay them to Earth, then go off to find its next sample. With the proper cameras, such a rover could gather information about the weathering rind, and about the makeup of many Mars' rocks.

The weathering rind mission should be done very soon. Such a mission could be done very economically. Two or three mini-rovers would be dropped to the planet's surface via parachutes. A small relay orbiter would be left in orbit. Each rover would find an appropriate sample, take the images, and send a signal to the orbiter. When the orbiter was in range, it would broadcast a ready to receive signal, and the rover could dump the images directly from the camera CCDs to the orbiter. This way, the rovers need no mass storage, and can get by with very simple electronics.

The information from these images is crucial in designing a proper rover for a detailed science sample return mission. If the weathering rind is several centimeters then a rover large enough to do rock coring is necessary for many experiments. If the rind is on the order of a millimeter, then a sample return mission using behavior-controlled mini-rovers (see [Miller90]) would be more than adequate, and much more cost-effective.

4. Conclusions

Controlling a rover through behavior-control can allow such a rover to be greatly reduced in size and complexity with little or no loss in capability. Planetary rover missions, in most cases, are rover missions because the mission objectives require the sensors to be in many different locations. All planetary missions are mass constrained. Behavior controlled mini-rovers are a way of getting beyond the mass constraints and increasing the effective mobility of the system. Mobility is increased by making the rovers more autonomous (detailed direction from Earth does not fit well with the behavior-control paradigm) and by being able to provide more rovers, dropped at more locations, for a given mass allotment, then is possible using traditional control or planning techniques. Behavior-control, when combined with nano-technolgy, also holds the promise of being able to undertake wholly new types of missions [Brooks89].

Acknowledgements: The research described in this paper was carried out by the Jet Propulsion Laboratory — California Institute of Technology under a contract with the National Aeronautics and Space Administration.

References

[Brooks86] Rodney A. Brooks, "A Robust Layered Control System for a Mobile Robot", IEEE Journal on Robotics and Automation, vol RA-2, no. 1, March 1986.

[Brooks89] Brooks, R.A., Flynn, A.M., Fast, Cheap, and out of Control: A Robot Invasion of the Solar System, *Journal of the British Interplanetary Society*, vol 42, #10, pp478-485, October 1989.

[Easter88] Easter, D.A., Buoni, C.M., McCauley, L.A., Mobility study for a Lunar rover, in the *Proceedings of the Mobile Robots III Conference*, SPIE Cambridge Symposiums, SPIE vol 1007, pp128-135, November 1988.

[Gat90] Gat, E., Slack, M.G., Miller, D.P., Firby, R.J., Path Planning and Execution Monitoring for a Planetary Rover, in *Proceeding of the IEEE International Conference on Robotics and Automation*, Cincinnati, OH, May 1990.

[Gat90b] Gat, E., Miller, D.P., *BDL: A Language for Programming Reactive Robotic Control Systems*, JPL Working paper, 1990.

[Miller89] Miller, D.P., Atkinson,D.J., Wilcox,B., Mishkin,A.H., Autonomous Navigation and Control of a Mars Rover, *Proceedings of the 11th IFAC Symposium on Automatic Control in Aerospace*, pp127-130, Tsukuba, Japan, July 1989.

[Miller90] Miller, D.P., Mini-rovers for Mars Exploration, in the *Proceedings of the Vision 21 Workshop*, NASA-Lewis Research Center, Cleveland, OH, April 1990.

[Pivirotto90] Pivirotto, D.S., Dias, W.C., *United States Planetary Rover Status - 1989*, NASA Jet Propulsion Laboratory Publication JPL 90-6, 1990.

[Ranger63] "Final Technical Report, Lunar Rough Landing Capsule Development Program," Aeronutronic Division Publication Number U-2007, February 1963.

[Waltman89] Waltman, S.B., Kaiser, W.J., Electron Tunnel Sensor Technology, *Journal of the British Interplanetary Society*, vol 42, #10, pp474-477, October 1989.